ISCR

Institute for Scientific Computing Research

Laboratory Directed Research and Development Project Research Summaries

CASC Center for Applied Scientific Computing

SAVAnTS: Scalable Algorithms for Visualization and Analysis of Terascale Science

Mark Duchaineau

Center for Applied Scientific Computing

Summary:

Wavelets and other hierarchies on large-scale scientific geometry form critically needed computational infrastructure to reduce the overwhelming data sizes and to make possible highly interactive exploration of scientific results. Several difficult challenges must be overcome to create such multi-resolution compression and display capabilities. First, appropriate grids must be devised on highly convoluted scientific surfaces, for which no methods have previously been known. Second, novel wavelet transforms are required that are completely local in their mathematical operations, so as to operate effectively on arbitrarily large data sets. Finally, display algorithms must be devised that read exactly the portion of the compressed data required for any given moment of interaction, and use this to feed graphics hardware and optimize its speed and accuracy. Overall, the goal is to achieve orders of magnitude improvements in storage sizes and interaction rates over current best practice.

During a given run, multi-physics simulation codes on LLNL supercomputers produce dozens of terabytes of data that form a vital component of the Stockpile Stewardship Management Program and other programs. Great strides are being made to increase the efficiency and accuracy of the codes by harnessing tens of thousands of processors using scalable algorithms. However, the efficient and accurate post-computation data handling and interactive exploration must also scale efficiently to reach LLNL's goal of a capability for productive 100 Tflop/s or greater simulation. This project introduces multi-resolution methods to data exploration activities that are especially critical to the Laboratory missions: the compact storage and fast display of variables on material boundaries, orthogonal cut planes of 3D field data, contour surfaces, and transparent volume renderings.

Our focus is on devising new wavelet transforms and other hierarchies, and applying them for compression and accelerated display of volumetric field data, material boundaries, and contour surfaces. We find that the useful information content of a 3D field such as pressure or density tends to be quite sparse. Because wavelets automatically find and exploit coherence in both space/time and frequency/scale, this sparse information content is readily compressed after the application of the appropriate wavelet transform.

Wavelets are well understood for regularly spaced grids filling a 3D block stored in a single computer's memory. However, innovations are required for large-scale Laboratory applications, including:

(1) highly adaptive or unstructured settings,

(2) arbitrary surfaces that typically can not be represented as anything resembling a regularly spaced grid, and

(3) data distributed over ten thousand processors.

Fundamentally, the work performed should be proportional to the sparse post-transform information content at as many stages as possible of the end-to-end data flow going from simulation to scientist. This leads to a suite of connected optimization problems that we have addressed.

Summary (continued):

Mark Duchaineau

Center for Applied Scientific Computing

During FY2001 we devised:

(1) a hierarchy-building process for surfaces, extended to exploit time coherence;

(2) new fast and accurate algorithms to extract material boundaries from volume fraction information and produce a related 3D fieldhierarchy;

(3) a shink-wrap surface remapping method that optimizes subsequent display-hierarchy accuracy;

(4) a volumetric-based surface compressor that overcomes the complex-topology limitations of explicit surface mappings;

(5) a spatial index remapper that automatically optimizes cache coherence and enables out-of-core computation on huge data sets;

(6) a magnifying-lens capability for hierarchical hardware-based transparent volume rendering;

(7) a wavelet compression technique that allows contour topology reservation;

(8) a memory-insensitive technique for surface simplification;

(9) a highly simplified and effective height-map display accelerator.

Several large-scale simulations were performed in the fall of 2000 on several thousand processors of the initial-delivery ASCI White supercomputer at LLNL. In collaboration with Farid Abraham of IBM Almaden Research, we have successfully studied for the first time the supersonic propagation of cracks and the formation of complex junction structures in metals. These unprecedented computations were made possible by our project's prototype software for wavelet compression and hierarchical display optimization. In one simulation on 5120 processors, it was demonstrated that it is possible to reduce the anticipated 25 terabytes of output to under a terabyte without perceptible degradation in visualization results, and that the compression work added only 10% to the simulation run time.

Overcoming the Memory Wall in SMP-Based Systems

Bronis R. de Supinski, Andy Yoo, Sally A. McKee, Frank Mueller, and Tushar Mohan

Center for Applied Scientific Computing

Summary:

Both CPU and memory speeds are increasing at exponential rates, as expressed in Moore's Law. Unfortunately, memory hardware is currently significantly slower than CPUs. For example, on snow, the ASCI White testbed system, an average of 87 floating point operations can be completed in the time required to load one operand from main memory. Even worse, CPU speeds are increasing faster than memory speeds; thus, the number of CPU cycles required to access memory is increasing. This divergence will exacerbate an existing problem for codes with large memory footprints, including the codes typically in use at LLNL: memory accesses dominate performance. Not only is the performance of many LLNL codes dominated by the cost of main memory accesses, but many current trends in computer architecture will lead to substantial degradation of the percentage of peak performance obtained by these codes. Many researchers anticipate a "Memory Wall" in which memory accesses imply an absolute performance limit, and improvements in CPU speed provide no performance benefit.

This project is extending dynamic access optimizations (DAO), a promising set of techniques for overcoming the Memory Wall, to symmetric multiprocessors (SMPs), which are common at LLNL. We expect future computer systems to become available to LLNL that use the novel techniques that we are designing to alleviate this problem in SMP-based systems. Further, our techniques complement other emerging mechanisms for improving memory system performance that will be the basis of future LLNL systems, such as processors-in-memory.

DAO techniques have shown significant promise to overcome the Memory Wall without requiring complex source code changes. These techniques change the order or apparent locations of memory accesses from those generated by the issuing program to ones that use the memory system more effectively without changing the results. For example, altering the execution order can exploit memory hardware characteristics such as interleaved memory banks and hot dynamic random access memory (DRAM) rows, while techniques that alter the apparent location can significantly increase cache hit ratios. DAO mechanisms can reduce run times of memory intensive portions of programs by factors of two to an order of magnitude. Previous projects investigating DAO focus on uniprocessor systems and require special-purpose hardware. Although promising, DAO techniques for uniprocessors do not target the systems in use at LLNL. All major LLNL computing resources are clusters with SMP nodes. Thus, we need DAO techniques that support simultaneous access to the memory system by multiple processors. Difficulties arise in SMPs for both types of DAO techniques, access reordering, and shadow memory. This project addresses these difficulties.

During FY01, the first year of the project, we established the infrastructure required for our research. The overall plan is not only to design DAO techniques for SMP-based systems, but to demonstrate that they significantly improve performance for typical LLNL codes. In order to achieve this, we will predict the benefits of these techniques with an analytic model that we are designing, and we will simulate the variety of techniques that we design, as well as implementing the most promising techniques on an actual SMP.

Summary (continued):

Bronis R. de Supinski, Andy Yoo, Sally A. McKee, Frank Mueller, and Tushar Mohan

Center for Applied Scientific Computing

In FY01, we have implemented a simulation that will serve as the basis of all of our simulation studies. In addition, we have gathered the input data of our analytic model for UMT, an LLNL 3D neutral particle transport code for unstructured meshes. This data includes both a characterization of the access patterns of UMT and detailed measurements of the memory systems of several SMP-based systems. All of this data was gathered with new tools that we implemented as part of this project.

An important aspect of our project is that we will not only design SMP-aware DAO techniques and demonstrate their efficacy through simulation, but we will also implement the techniques on an actual system and demonstrate their benefit for UMT on that system. Thus, two of our FY01 milestones involved the selection of an experimental system and an evaluation on the hardware we would use to implement our DAO techniques. After the project started, we established a close collaboration with SRC Computers of Colorado Springs, which includes privileged use of their new SMP. Since its memory controllers, cache coherence controllers, network bridges, and switches are all implemented with field programmable gate arrays (FPGAs), we can implement our SMP-aware DAO techniques directly in the hardware by changing the FPGA programs of these devices.

In FY02, the second year of this project, we will simulate and implement the first SMP-aware DAO techniques and refine the techniques based on our initial results. We will investigate the performance of SMP-aware DAO techniques in the presence of message passing memory traffic and explore techniques that target SOC and PIM in the third year of the project.

Sapphire: Scalable Pattern Recognition for Large-Scale Scientific Data Mining

Chandrika Kamath

Center for Applied Scientific Computing

Summary:

There is a widening gap between our ability to collect data and our ability to analyze it. This problem of data overload has become a serious impediment to scientific advancement in areas as diverse as counter-proliferation, the Advanced Simulation and Computing (ASCI) department, astrophysics, computer security, and climate modeling. To improve the way in which scientists extract information from their data, we are developing a new generation of tools and techniques based on data mining.

Data mining is the semi-automated discovery of patterns, associations, anomalies, and statistically significant structures in data. In the first step of data preprocessing, high-level features are extracted from the data; in the second step of pattern recognition, the features are used to identify and characterize patterns in the data. In this project, we have developed scalable algorithms for the pattern recognition task of classification. We have improved their performance, without sacrificing accuracy. We have demonstrated these techniques using an astronomy application, namely the detection of radio-emitting galaxies with a bent-double morphology in the FIRST survey.

In FY2001, we focused on three tasks: (a) improving the performance of decision tree algorithms, (b) identifying bent-double galaxies in the FIRST survey, and (c) incorporating our research into software to make it easily accessible to LLNL scientists. In decision trees, we focused on ensembles of trees, where the results of several trees are combined through simple voting. We invented two new ways of creating ensembles by randomizing the decision at each node of the tree. The first approach uses a random sample of the instances for each feature. The second approach uses histograms and randomly selects a split point in an interval around the best bin boundary. Using public-domain data sets, we showed that both techniques were more accurate than single trees and competitive in accuracy with other techniques for creating ensembles, but faster.

For the bent-double problem, we focused on galaxies composed of three blobs. Using principal component analysis and exploratory data analysis techniques, we first identified the key features discriminating bent-doubles from non-bent-doubles. This reduced the number of features from 103 to 31. We next input these features to our decision tree software as well as the generalized linear model software from S-PLUS. Varying the number of input features, we created three models for each of the two classifiers. These models were used to classify unseen galaxies and the results were communicated to our collaborators on the FIRST project. Galaxies where all six models agreed were considered bent-doubles with high probability, while those with some disagreement resulted in a lower probability.

We completed the Beta version of our software in December 2000 and Version 1.0.0 in September 2001. This includes the recent algorithms developed for ensembles and the evolutionary algorithm-based oblique decision trees that we had developed last year. We also enhanced the decision trees with several pruning options, splitting criteria, and split finders.

Summary (continued):

Chandrika Kamath
Center for Applied Scientific Computing

We co-edited one book, published 17 papers in conferences, journals, and books, and presented our work at 15 conferences and workshops. We also filed three new records of invention and three patent applications based on earlier work in FY2000. We co-organized two workshops and one week-long program at the NSF Institute for Pure and Applied Mathematics at UCLA, gave three tutorials at conferences, and actively participated in university collaborations.